arXiv:1602.01522v1 [stat.ME] 4 Feb 2016

Risk estimation for high-dimensional lasso regression

Darren Homrighausen
Department of Statistics
Colorado State University
darrenho@stat.colostate.edu

Daniel J. McDonald
Department of Statistics
Indiana University, Bloomington
dajmcdon@indiana.edu

Version: February 5, 2016

Abstract

In high-dimensional estimation, analysts are faced with more parameters p than available observations n, and asymptotic analysis of performance allows the ratio $p/n \to \infty$. This situation makes regularization both necessary and desirable in order for estimators to possess theoretical guarantees. However, the amount of regularization, often determined by one or more tuning parameters, is integral to achieving good performance. In practice, choosing the tuning parameter is done through resampling methods (e.g. cross-validation), generalized information criteria, or reformulating the optimization problem (e.g. square-root lasso or scaled sparse regression). Each of these techniques comes with varying levels of theoretical guarantee for the low- or high-dimensional regimes. However, there are some notable deficiencies in the literature. The theory, and sometimes practice, of many methods relies on either the knowledge or estimation of the variance parameter, which is difficult to estimate in high dimensions. In this paper, we provide theoretical intuition suggesting that some previously proposed approaches based on information criteria work poorly in high dimensions. We introduce a suite of new risk estimators leveraging the burgeoning literature on high-dimensional variance estimation. Finally, we compare our proposal to many existing methods for choosing the tuning parameters for lasso regression by providing an extensive simulation to examine their finite sample performance. We find that our new estimators perform quite well, often better than the existing approaches across a wide range of simulation conditions and evaluation criteria.

Keywords: Model selection; tuning parameter selection; prediction; variance estimation

1 Introduction

Suppose we are given a data set, $Z_1, \ldots, Z_n$, of paired observations including a covariate $X_i \in \mathbb{R}^p$ and its associated response $Y_i \in \mathbb{R}$ such that $Z_i^\top = (X_i^\top, Y_i)$. Concatenating the covariates row-wise into a matrix, we obtain the design matrix $\mathbb{X} = [X_1, \ldots, X_n]^\top \in \mathbb{R}^{n \times p}$.

We assume that the relationship between the covariate and response is of the form

$$Y = \mathbb{X}\beta^* + \epsilon, \tag{1}$$

where $\epsilon \sim (0, \sigma^2 I)$, meaning the entries of ϵ are mean zero with uncorrelated components each having variance σ^2.

When $p > n$, estimation of the linear model requires some structural assumptions on the properties of β^* resulting in a penalized linear model. We will focus mainly on the lasso with

$$\widehat{\beta}(\lambda) = \underset{\beta \in \mathbb{R}^p}{\operatorname{argmin}} \frac{1}{n} ||Y - \mathbb{X}\beta||_2^2 + \lambda ||\beta||_1 , \tag{2}$$

where $\lambda \geq 0$ is a tuning parameter and $||\cdot||_2$ and $||\cdot||_1$ are the ℓ_2 (Euclidean) and ℓ_1 norms respectively. Similar M-estimators with different penalties include, among others, ridge regression, the group lasso [29], and the smoothly clipped absolute deviation penalty [SCAD, 8]. We concentrate on the lasso, though we will occasionally also reference ridge regression,

$$\widehat{\beta}_{ridge}(\lambda) = \underset{\beta \in \mathbb{R}^p}{\operatorname{argmin}} \frac{1}{n} ||Y - \mathbb{X}\beta||_2^2 + \lambda ||\beta||_2^2 = (\mathbb{X}^\top \mathbb{X} + \lambda I_p)^{-1} \mathbb{X}^\top Y,$$

because it has a closed form which can provide intuition.

For ridge regression, a unique solution always exists for $\lambda > 0$, although, for λ small enough, numerical issues may intercede. For lasso, by convexity there is always at least one solution to equation (2), although if $\operatorname{rank}(\mathbb{X}) = n$, there may be multiple minimizers [see 24, for details]. In this case, we refer to 'the' solution as the outcome of the particular minimization technique used (e.g. LARS [7]). We will also consider some modifications to equation (2) which attempt to eliminate the influence of tuning parameters through innovative optimization routines (see Section 4.2 for a more detailed description).

The theoretical optimality properties that exist in the literature for penalized regression rely on appropriate tuning parameter selection. Under restrictions on the design matrix $\mathbb{X}$, the distribution of ϵ, and the sparsity pattern of β^*, Bunea et al. [4] show that, as long as the number of nonzero entries in β^* does not increase too quickly, deviations of the prediction risk larger than $\sigma^2 \log(p)/n$ go to zero if $\lambda_n = a\sigma\sqrt{\log(p)/n}$ for some constant a. Likewise, deviations in the distance between $\widehat{\beta}(\lambda)$ and β^* of order $\sigma\sqrt{\log(p)/n}$ have vanishing probability. While theoretical results of this type provide comfort that a data analyst's procedure will eventually perform well given sufficient data, they depend on unknown quantities such as σ^2, the noise distribution, and other constants.

In practice, many methods for empirically choosing λ have been proposed. The aim of this paper is to review these techniques by (1) introducing a suite of novel risk estimation methods that are straightforward to compute and perform well empirically, (2) contrasting these new risk estimation methods with existing, superficially similar GIC-based methods, and, lastly, (3) running a comprehensive simulation study over a broad range of data generating scenarios. This investigation reveals some existing deficiencies in current high-dimensional risk estimation approaches and highlights fruitful research directions, particularly the relationship between risk estimation and variance estimation in high dimensions while justifying our proposal.

Notation: For any vector β, we denote $\mathcal{S} = \mathcal{S}(\beta) := \{j : \beta_j \neq 0\}$, and $\mathbb{X}_{\mathcal{S}(\beta)}$ to be the columns of the design matrix selected by β. We write $\mathcal{S}^* = \mathcal{S}(\beta^*)$, and $s^* = |S^*|$.

2 Risk estimation in high-dimensional regression problems

Under the model in equation (1), the ℓ_2 prediction risk of a coefficient vector β can be written [20]

$$\begin{aligned} n^{-1}\mathbb{E}\,||\mathbb{X}\beta - \mathbb{X}\beta^*||_2^2 &= n^{-1}\mathbb{E}\,||\mathbb{X}\beta - Y||_2^2 - \sigma^2 + 2n^{-1}\sum_{i=1}^{n}\mathrm{Cov}(\widehat{Y}_i, Y_i), \\ &= n^{-1}\mathbb{E}\,||\mathbb{X}\beta - Y||_2^2 - \sigma^2 + 2n^{-1}\sigma^2\mathrm{df}, \end{aligned}$$

where the *degrees of freedom* [6] of the prediction $\widehat{Y} = \mathbb{X}\beta \in \mathbb{R}^n$ is defined to be

$$\mathrm{df} = \frac{1}{\sigma^2}\sum_{i=1}^{n}\mathrm{Cov}(\widehat{Y}_i, Y_i),$$

with $\mathrm{Cov}(\widehat{Y}_i, Y_i) = \mathbb{E}\left[(\widehat{Y}_i - \mathbb{E}\widehat{Y}_i)(Y_i - \mathbb{E}Y_i)\right]$. Therefore, a suite of sensible estimators of the ℓ_2 prediction risk is produced via

$$\widehat{R}_\beta(\widehat{\sigma}^2, C_n) = n^{-1}\,||\mathbb{X}\beta - Y||_2^2 - \widehat{\sigma}^2 + C_n\widehat{\sigma}^2\widehat{\mathrm{df}}, \tag{3}$$

where $\widehat{\sigma}^2$ is an estimator of σ^2, C_n is a function depending on n, and $\widehat{\mathrm{df}}$ is an estimator of the degress of freedom. For simplicity, we occasionally omit any arguments to $\widehat{R}$ that aren't directly relevant to the discussion at hand. Additionally, we define $\widehat{\mathrm{train}}_\beta \equiv \widehat{R}_\beta(0,0)$ to be the training error of β and write $\widehat{R}_\lambda \equiv \widehat{R}_{\beta(\lambda)}$ when β is indexed by the tuning parameter λ.

The choice $C_n = 2n^{-1}$, along with unbiased estimators $\widehat{\mathrm{df}}$ and $\widehat{\sigma}^2$, makes $\widehat{R}_\beta$ an unbiased estimator of $n^{-1}\mathbb{E}\,||\mathbb{X}\beta - \mathbb{X}\beta^*||_2^2$. In the low dimensional case, where $\widehat{\beta}(0)$ is the least squares solution and $\widehat{\sigma}^2 = (n-p)^{-1}\left|\left|Y - \mathbb{X}\widehat{\beta}(0)\right|\right|_2^2$, $\widehat{R}_{\widehat{\beta}(0)}(\widehat{\sigma}^2)$ is the classical Mallow's Cp [17]. However, when $\mathrm{rank}(\mathbb{X}) = n$, the choice of an estimator of the noise variance σ^2 is far from straightforward, and different choices of the penalty, C_n, provide different risk estimators (e.g. AIC or BIC). The lasso path algorithm `R` package `lars` provides a Cp-like score for choosing the model. However, the `lars` Cp statistic is undefined when $\mathrm{rank}(\mathbb{X}) = n$ as it uses the classical residual-sum-of-squares variance estimator.

If a predictor is linear in Y, that is $\widehat{Y} = HY$ for some matrix H, then $\mathrm{df} = tr(H)$, i.e. the trace, or sum of the diagonal entries. Hence, for $\widehat{\beta}(0)$, $\mathrm{df} = \mathrm{rank}(\mathbb{X})$, which equals p if $\mathbb{X}$ has linearly independent columns. However, in the case of lasso, df must be estimated (hence the notation $\widehat{\mathrm{df}}$) as lasso is not linear in the response. In [25, 32] it is shown that for the lasso, the degrees of freedom of $\widehat{Y} = \mathbb{X}\widehat{\beta}(\lambda)$ is equal to $\mathbb{E}[\mathrm{rank}(\mathbb{X}_{\mathcal{S}(\lambda)})]$, suggesting the natural unbiased estimator $\widehat{\mathrm{df}} = \widehat{\mathrm{df}}(\lambda) = \mathrm{rank}(\mathbb{X}_{\mathcal{S}(\lambda)})$. As the high dimensional setting precludes the standard variance estimator $\widehat{\sigma}^2$, the Stein framework requires an alternative. Next, we discuss some high-dimensional variance estimators that have been recently advocated.

2.1 Variance estimation

The literature on variance estimation in high dimensions is a quickly growing field. In this paper, we explore three high-dimensional variance estimation approaches.

The first two approaches start by finding the solution path for the lasso problem (that is, the solutions to equation (2) at the values of λ which result in a change in $\mathcal{S}$) and choosing $\widehat{\beta}(\widehat{\lambda}_{CV})$ by minimizing a K-fold cross-validation estimator of the risk (see Section 4.1 for the details of cross validation). With this coefficient estimate, the squared ℓ_2 norm of the residuals can be used as a variance estimate, i.e.

$$\widehat{\sigma}^2_{CV} = \frac{1}{n - \widehat{\mathrm{df}}(\widehat{\lambda}_{CV})} \left|\left|Y - \mathbb{X}\widehat{\beta}(\widehat{\lambda}_{CV})\right|\right|_2^2. \tag{4}$$

Alternatively, a restricted maximum likelihood-type method can be formed by examining the orthogonal complement of the projection onto the column space of $\mathbb{X}_{\mathcal{S}(\widehat{\lambda}_{CV})}$: $H^{\perp}_{CV}$. Using this projection we define

$$\widehat{\sigma}^2_{RMLE} = \frac{1}{tr(H^{\perp}_{CV})} \left|\left|H^{\perp}_{CV}Y\right|\right|_2^2.$$

For the lasso, $\mathrm{trace}(H^{\perp}_{CV}) = \mathrm{trace}(I - H_{CV}) = n - \mathrm{rank}\left(\mathbb{X}_{\mathcal{S}(\widehat{\lambda}_{CV})}\right)$, which implies $\mathrm{trace}(H^{\perp}_{CV}) = n - \widehat{\mathrm{df}}(\widehat{\lambda}_{CV})$. Hence these two variance estimators differ only in the size of the residuals.

The third variance estimation method we consider is known as refitted cross-validation (RCV) [9]. After randomly splitting the data in half, $\mathbb{X}_{\mathcal{S}(\widehat{\lambda}_{CV})}$ is formed on the first half and $\widehat{\sigma}^2_1$ is formed via equation (4), using the Y and $\mathbb{X}$ values from the second half. The procedure is then repeated, exchanging the roles of the halves, producing $\widehat{\sigma}^2_2$. A final estimate is formed via $\widehat{\sigma}^2_{RCV} = (\widehat{\sigma}^2_1 + \widehat{\sigma}^2_2)/2$.

In a comprehensive simulation study, Reid et al. [18] find that $\widehat{\sigma}^2_{CV}$ is the most reliable estimator out of those cited above, although, as pointed out in Fan et al. [9], it appears to have a downward bias relative to $\widehat{\sigma}^2_{RCV}$. However, this doesn't mean that any of the above methods will necessarily produce superior performance as a plug-in variance estimator for risk estimation. In fact, due to the nature of projections, for the lasso

$$\left|\left|H^{\perp}_{CV}Y\right|\right|_2^2 \leq \left|\left|Y - \mathbb{X}\widehat{\beta}(\widehat{\lambda}_{CV})\right|\right|_2^2,$$

thus it must hold that $\widehat{\sigma}^2_{RMLE} \leq \widehat{\sigma}^2_{CV}$. Hence, $\widehat{R}(\widehat{\sigma}^2_{RMLE})$ penalizes model complexity less than $\widehat{R}(\widehat{\sigma}^2_{CV})$. In Section 5, our simulations show that, when choosing $C_n = 2n^{-1}$, $\widehat{R}(\widehat{\sigma}^2_{CV})$ results in lower prediction risk, better estimation consistency, and higher precision, while $\widehat{R}(\widehat{\sigma}^2_{RMLE})$ has better recall.

Armed with any of the above high-dimensional variance estimators, $\widehat{\sigma}^2_{HD}$, we can form an estimator of β^* via

$$\widehat{\beta} = \underset{\beta}{\mathrm{argmin}}\, \widehat{R}_{\beta}(\widehat{\sigma}^2_{HD}, C_n).$$

Note that in regression problems, there is a long history [1, 19] of using (estimates of) information criteria for model selection. Though coming from very different theoretical derivations, unbiased risk estimation and generalized information criteria (GIC) can produce

very similar procedures (especially in low-dimensional situations) and hence are sometimes treated equivalently. However, in the next section we demonstrate that constructing GIC in high dimensions requires care.

3 Generalized information criteria

The risk estimator $\widehat{R}_\beta(\widehat{\sigma}^2)$ in equation (3) shares many similarities with generalized information criteria under Gaussian noise. However, as we demonstrate below, a key component to defining GIC in high dimensions is knowledge of the scale parameter. We discuss GIC in some detail here due to its frequent use in tuning parameter selection.

There are two commonly proposed GIC forms for linear models of the form in equation (1) depending on whether or not σ^2 is known (or assumed known). If σ^2 is unknown and ϵ is Gaussian, then GIC takes the form

$$\text{info}(C_n, g) := \log\left(\widehat{\text{train}}_\beta\right) + C_n\, g(\widehat{\text{df}}), \tag{5}$$

where C_n is a constant that depends only on n, and $g : [0, \infty) \to \mathbb{R}$ is a fixed function. In the context of choosing the tuning parameter λ in the lasso problem, this GIC form is frequently suggested in the literature [for example 3, 8, 12, 23, 26]. The specific choices $g(x) = x$ and $C_n = 2/n$ or $C_n = \log(n)/n$ are commonly referred to as AIC and BIC, respectively. Additionally, generalized cross validation (GCV) is defined to be

$$\text{GCV} = \frac{\widehat{\text{train}}_\beta}{(1 - \widehat{\text{df}}/n)^2}$$

If we write this on the log scale, it takes the form of equation (5) with $g(x) = \log(1 - x/n)$ and $C_n = -2$.

Alternatively, if the scale parameter is considered known [10, 31] GIC matches $\widehat{R}_\beta(\sigma^2)$, where the true, known variance is used instead of an estimator. However, the σ^2 known assumption, though common in theoretical derivations, is only really appropriate when $p/n \to 0$ and the classical residual sums of squares estimator of the variance can be used. High-dimensional variance estimation is still a developing topic, and its properties as a plug-in estimator for model selection have not been investigated.

Classical asymptotic arguments underlying information criteria apply only for p fixed and rely on maximum likelihood estimates (or Bayesian posteriors) for all parameters including the noise. This theory breaks down in the high dimensional setting for two reasons. First, if $\text{rank}(\mathbb{X}) = p$ but is still allowed to increase, recent work has developed new information criteria with supporting asymptotic results. For example, the criterion developed in Wang et al. [27] selects the correct model asymptotically even if $p \to \infty$ as long as $p/n \to 0$. Additionally, Flynn et al. [12] investigate a variety of properties and GIC-methods, but only when $\text{rank}(\mathbb{X}) = p$. If p is allowed to increase more quickly than n, theoretical results assume σ^2 is known to get around the difficult task of high-dimensional variance estimation [5, 10, 16, 30].

Though there is this substantial literature advocating each GIC form, for the high-dimensional setting there are several notable gaps. First and most seriously, $\text{info}(C_n, g)$

from equation (5) is unusable if $\text{rank}(\mathbb{X}) = n$. This is because it is always possible to achieve $\widehat{\text{train}}_\beta = 0$ by allowing $\lambda = 0$. Therefore, $\text{info}(C_n, g)$ must approach $-\infty$ unless g increases to ∞ faster. In many cases, these criteria will thus choose a saturated model. Nonetheless, this approach is frequently recommended [3, for example].

To reinforce these points, we provide two simple examples demonstrating the empirical faillures of $\text{info}(C_n, g)$ and $\widehat{R}_\beta(\sigma^2)$, the risk estimate with known variance. The first shows that when $\text{rank}(\mathbb{X}) < p$, GIC will always select p covariates. The second shows that assuming the variance is known when it is not leads to unstable model selection. Despite the intuitive nature of these claims, both situations appear in recent literature.

Example 1

Consider the following regression dataset. Let

$$Y = \frac{\sigma}{\sqrt{2}} \begin{bmatrix} 1 \\ -1 \end{bmatrix}, \qquad \text{and} \qquad \mathbb{X} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & -1 & 1 \\ -1 & 1 & -1 \end{bmatrix}.$$

In this no noise case, Y is a scalar multiple of each column of X. This setup could easily be obtained with two independent draws from the population model $y_i = x_{i1} + \epsilon_i$ where $\epsilon_i \sim (0, \sigma^2)$ and subsequently normalizing $\mathbb{X}$ and Y. Repeating this experiment numerous times will lead to datasets which look like this one on average.

As ridge regression has a closed form-solution, we first look at its GIC properties for intuition. One can show that

$$\text{df}(\lambda) = \frac{3}{\lambda + 3} \qquad \text{and} \qquad \widehat{R}_\lambda = \frac{\sigma^2 \lambda^2}{(\lambda + 3)^2},$$

and so,

$$\text{info}(C_n, g) = \log\left(\frac{\sigma^2 \lambda^2}{2(\lambda + 3)^2}\right) + C_n \, g\left(\frac{3}{\lambda + 3}\right).$$

For $0 < \lambda < 1$, $\sigma^2\lambda^2/16 \leq \widehat{R}_\lambda \leq \sigma^2\lambda^2/9$, so $\log(\widehat{R}_\lambda) \to -\infty$ like $\log(\lambda)$ as $\lambda \to 0$. Hence, minimizing $\text{info}(C_n, g)$ will choose $\lambda = 0$ unless the second term increases at least as fast as $-\log(\lambda)$, that is we require constants c and C such that $g\left(\frac{3}{\lambda+3}\right) \geq C \log(1/\lambda)$ for all $\lambda < c$. We see immediately that AIC and BIC (which both have $g(x) \equiv x$) will always fail this test and must select $\lambda = 0$. For GCV, the issue is a bit more subtle. In this example, as $1 = \text{rank}(\mathbb{X}) < n = 3$, $\log(n - \text{df}) \to 0$ and hence GCV selects $\lambda = 0$ as well. If $\text{rank}(\mathbb{X}) = n$, then $\log(n - \text{df}) \to -\infty$ and hence the rate that $\log(n - \text{df})$ goes to $-\infty$, along with magnitude of the constants involved, determine whether the trivial $\lambda = 0$ solution is returned.

In Figure 1, we plot AIC for $\lambda \in [1 \times 10^{-5}, 1]$ for both ridge (left plot) and lasso (right plot) regression. Using AIC would have us report the unregularized model; that is picking the smallest λ which we allow and using a least squares solution. Of course had we simply used one column alone with a standard linear model, we would have both fit the data perfectly and precisely estimated the single parameter.

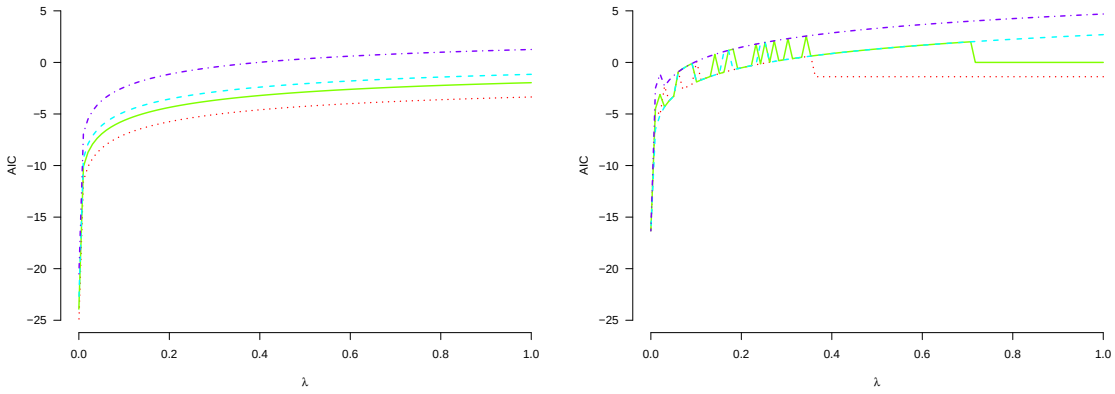

Figure 1: The left plot shows AIC as we vary λ from 1×10^{-5} to 1 for the small numerical example. The right plot shows the same setup but using lasso instead. The different curves have σ set to .5 (red, dotted), 1 (green, solid), 1.5 (cyan, dashed) and 5 (purple, dot dash). Notice that we would always choose the unregularized, $\lambda = 0$ model.

Example 2

Addtionally, we perform a small simulation to explore $\widehat{R}(\sigma^2)$. When this criterion is used in the literature, σ^2 is usually considered known for theoretical purposes [10] or set to a fixed value [28, simulation study].

We generate draws according to the model in equation (1), such that $n = 30$, $p = 150$, and β^* has one nonzero coefficient drawn from the standard Laplace distribution. In Figure 2, we explore four methods for choosing λ for the lasso. Clockwise from top left these methods are the fixed and "known" σ approach given by $\widehat{R}_\lambda(\sigma^2 = 1)$, $\widehat{R}_\lambda(\widehat{\sigma}^2_{\text{CV}})$, $\widehat{R}_\lambda(\widehat{\sigma}^2_{\text{RCV}})$ (see Section 2.1 for definitions of these variance estimators), and lastly info$(C_n = 2, g(x) = x)$, which corresponds to AIC.

As expected, $\widehat{R}(\sigma^2 = 1)$ performs quite poorly when σ is far from 1. Also, info$(C_n = 2, g(x) = x)$ continues to choose the unregularized solution, as predicted by the previous example. The other two, $\widehat{R}(\widehat{\sigma}^2_{\text{CV}})$ and $\widehat{R}(\widehat{\sigma}^2_{\text{RCV}})$, perform much better.

4 Additional tuning parameter selection methods

There are many procedures for choosing a particular lasso solution. Here, we discuss cross-validation, two modern methods, both of which propose an optimization problem related to that for lasso but designed to eliminate the need for tuning parameters, and a two-stage method which was motivated by our simulations. We examine each of these approaches empirically in Section 5.

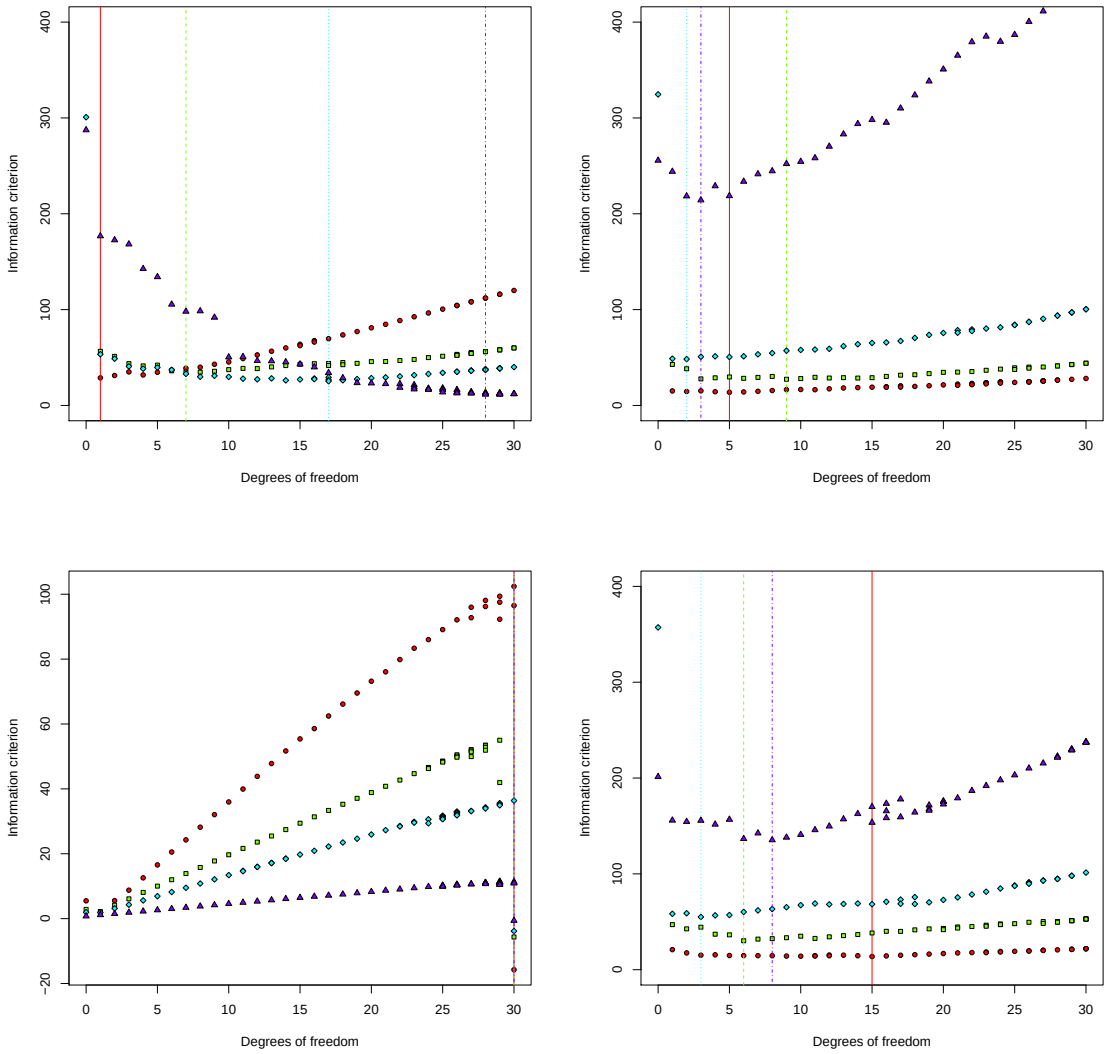

Figure 2: We use four different values for σ: $\sigma = 0.5$ (red, solid, circles), $\sigma = 1$ (green, dashed, squares), $\sigma = 1.5$ (cyan, dotted, diamonds), $\sigma = 5$ (violet, dash-dot, triangles). Risk estimation methods, clockwise from top left: $\widehat{R}(\sigma^2 = 1)$, $\widehat{R}(\widehat{\sigma}^2_{\mathrm{CV}})$, $\widehat{R}(\widehat{\sigma}^2_{\mathrm{RCV}})$, and info$(C_n = 2, g(x) = x)$. Notice that info$(C_n = 2, g(x) = x)$ always selects the unregularized model and $\widehat{R}(\sigma^2 = 1)$ depends significantly on σ.

4.1 Cross-validation

Frequently [for example 14, 15, 32], the recommended technique for selecting λ is through *K-fold cross-validation* (CV). Letting $V_n = \{v_1, \ldots, v_K\}$ be a partition of $\{1, \ldots, n\}$, then

$$CV(\lambda; V_n) := \frac{1}{K} \sum_{v \in V_n} \frac{1}{|v|} \sum_{r \in v} \left(Y_r - X_r^\top \widehat{\beta}^{(v)}(\lambda) \right)^2,$$

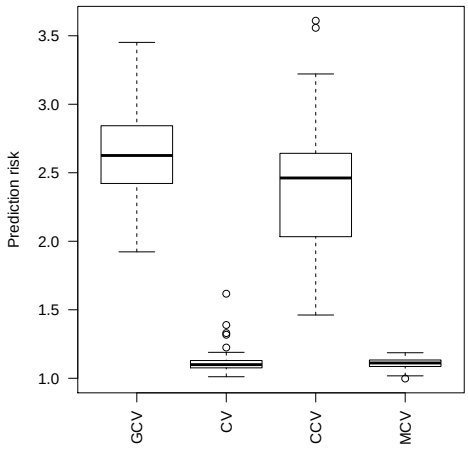

Figure 3: A comparison of cross-validation methods. Note that while GCV is more related to GIC-based methods, we include it here due to its historical association with cross-validation.

where $\widehat{\beta}^{(v)}(\lambda)$ is the lasso estimator in equation (2) with the observations in the validation set v removed, and $|v|$ indicates the cardinality of the set v. We define $\widehat{\lambda}_{CV} = \operatorname{argmin} CV(\lambda; V_n)$. We will use $K = 10$ as this is a popular choice in practice, and it has performed well in our experience.

Several adaptations for cross-validation have recently been suggested. We include two such methods here. The first, Modified Cross-Validation [MCV, 28], seeks to correct for a bias in CV induced by the lasso penalty. The second method, Consistent Cross-Validation [CCV, 11], finds the sequence of selected models by varying λ and uses cross-validation to select amongst that sequence.

For this paper, we use `glmnet` to fit CV and MCV and use `lars` to find the sequence of models for CCV and to compute the lasso path for GCV. Based on our simulation results, GCV, and CCV all dramatically underperform CV and MCV (See Figure 3 for a plot of typical results, though we defer a precise definition of prediction risk and a discussion on the particulars of the simulation conditions to Section 5). Therefore, for simplicity we will not include either GCV or CCV in subsequent comparisons.

4.2 Modified lasso-type methods

Alternatives to choosing the lasso tuning parameter without minimizing a risk estimate have recently been introduced. These methods claim to be 'tuning parameter free', which in theory means that the tuning parameters can be set independent of the data.

First, Sun and Zhang [21] develop 'scaled sparse regression' (SSR), which uses the fact that the optimal choice of λ_n for lasso is asymptotically proportional to σ.

By recasting the lasso problem as

$$\widehat{\beta}_{SSR} = \underset{\beta,\sigma}{\operatorname{argmin}} \frac{1}{2n\sigma} ||Y - \mathbb{X}\beta||_2^2 + \frac{(1-a)\sigma}{2} + \lambda_0 ||\beta||_1 ,$$

and fixing λ_0 and a, the authors develop theory for tuning-free lasso with simultaneous variance estimation. Though this is a promising approach, the objective function is not convex, hence the variance and the lasso solution are iteratively computed and the solutions tend to be sensitive to the starting values. Nonetheless, SSR enjoys attractive theoretical properties and has demonstrated some notable empirical successes.

Alternatively, Belloni et al. [2] suggest the $\sqrt{\text{lasso}}$ as a modification of the lasso problem

$$\widehat{\beta}_{\sqrt{\text{lasso}}} = \underset{\beta}{\operatorname{argmin}} \frac{1}{\sqrt{n}} ||Y - \mathbb{X}\beta||_2 + \frac{\lambda_n}{n} ||\beta||_1 . \tag{6}$$

Appealing to asymptotic arguments, Belloni et al. [2] show that the minimizer of equation (6) achieves near oracle performance if $\lambda_n = c\sqrt{n}\Phi^{-1}(1 - \alpha/(2p))$, which does not depend on σ. Here, Φ^{-1} is the quantile function for the standard Gaussian distribution. As $\sqrt{\text{lasso}}$ tended to pick the correct model but with overly regularized coefficient estimates, we will additionally examine a refitted version of $\sqrt{\text{lasso}}$ in which the unregularized least squares solution of Y on $\mathbb{X}_{\mathcal{S}(\widehat{\beta}_{\sqrt{\text{lasso}}})}$ is reported.

4.3 Two-stage method

Our experiments show that GCV tends to dramatically under regularize. Likewise, setting $C_n = n^{-1}\log(n)$ in $\widehat{R}_\beta(\widehat{\sigma}^2, C_n)$ tends to over regularize. Hence, we investigate a two-stage method whereby an intial screening is performed by selecting $\widehat{\lambda}_{GCV}$ and forming $\mathcal{S}_{\widehat{\lambda}_{GCV}}$. Then, we minimize $\widehat{R}_\beta(\widehat{\sigma}^2, C_n = n^{-1}\log(n))$ to produce a final solution.

5 Simulations

In the remainder of this paper, we will evaluate many of the proposed model selection methods in the high-dimensional setting using several criteria.

5.1 Simulation parameters

For these simulations, we consider a wide range of possible conditions by varying the correlation in the design, ρ; the number of parameters, p; the sparsity, α; and the signal-to-noise ratio, SNR. In all cases, we let $n = 100$ and set $\sigma^2 = 1$.

The design matrices are produced in two steps. First, $X_{ij} \overset{i.i.d.}{\sim} N(0,1)$ for $1 \leq i \leq n$ and $1 \leq j \leq p$, forming the matrix $\mathbb{X} \in \mathbb{R}^{n \times p}$. Second, correlations are introduced by defining a matrix D that has all off-diagonal elements equal to ρ and diagonal elements equal to one. Then, we redefine $\mathbb{X} \leftarrow \mathbb{X}D^{1/2}$. For these simulations, we consider correlations $\rho = 0.1,\ 0.5,\ 0.8$.

Abbreviation	Method
CV-10-Fold	10-fold cross validation
MCV	Modified Cross Validation
R-CV-2	$\widehat{R}_\beta(\widehat{\sigma}^2_{CV}, C_n = 2/n)$
R-RMLE-2	$\widehat{R}_\beta(\widehat{\sigma}^2_{RMLE}, C_n = 2/n)$
R-RCV-2	$\widehat{R}_\beta(\widehat{\sigma}^2_{RCV}, C_n = 2/n)$
R-CV-logn	$\widehat{R}_\beta(\widehat{\sigma}^2_{CV}, C_n = \log(n)/n)$
2-stage	Two-stage "GCV" discussed in Section 4.3
SSR	Scaled sparse regression
SQRT	$\sqrt{\text{lasso}}$
SQRT refitted	OLS estimation on the model selected with $\sqrt{\text{lasso}}$

Table 1: List of methods and abbreviations used in our empirical study

For sparsity, we define $s^* = \lfloor n^\alpha \rfloor$ and generate the s^* non-zero elements of β^* from a Laplace distribution with parameter 1. We let α be 0.1, 0.4 and 0.7, which corresponds to 1, 6, and 25 non-zero elements respectively. We vary the signal-to-noise ratio, defined to be $\text{SNR} = \beta^{*\top}\mathbb{E}[\mathbb{X}\mathbb{X}^\top]\beta^*/\sigma^2 = \beta^{*\top}\mathbb{E}[\mathbb{X}\mathbb{X}^\top]\beta^*$ from $\text{SNR} = 0.1, 1.0, 10$. Note that as SNR increases the observations go from a high-noise and low-signal regime to a low-noise and high-signal one. We let $p = 200$, 500, and 1000.

Lastly, we consider two different noise distributions, $\epsilon_i \sim N(0,1)$ and $\epsilon_i \sim 3^{-1/2}t(3)$. Here $t(3)$ indicates a t distribution with 3 degrees of freedom and the $3^{-1/2}$ term makes the variance equal to 1. As the results for these noise distributions are quite similar, we only present the Gaussian simulations. Furthermore, while we have simulated all combinations of these parameters and distributions, we include only a subset here for brevity.

5.2 Implementation of methods and notation

For ease of reference, Table 1 displays all of the risk estimators for which we present simulations. Since some of these methods rely on numerical optimization routines, it is important to discuss the particular implementation of the solvers used to generate $\widehat{\beta}(\lambda)$. Two widely used implementations for lasso are `glmnet` [13], which uses coordinate descent, and `lars`, which leverages the piece-wise linearity of the lasso solution path. The package `glmnet` is much faster than `lars`, however, `glmnet` only examines a grid of λ values and returns an approximate solution at each λ (due to the iterative nature of the algorithm). Additionally, `glmnet` suffers from numerical stability issues for small λ values when $p > n$.

Because it uses a grid of λ values, `glmnet` works better than `lars` for CV-10-Fold and to find $\widehat{\sigma}^2_{CV}$, $\widehat{\sigma}^2_{RCV}$, and $\widehat{\sigma}^2_{RMLE}$. Once we have these estimates, we use `lars` to find the entire lasso solution path and choose the minimizer of $\widehat{R}(\widehat{\sigma}^2)$ along the path as the final model. The reason for this approach is that the `lars` path will necessarily change for different cross-validation folds, and hence, the values of λ will not be comparable. Using the grid from `glmnet` allows us to calculate the CV estimates at the same values of λ across all folds. Then we can use `lars` to get the final estimate $\widehat{\beta}$. As for optimizing the modified lasso problems' objective functions, we use the `R` package `scalreg` to fit SSR and the `R` package

`flare` to fit the $\sqrt{\text{lasso}}$. For `scalreg`, we choose the starting point for the iteration via the quantile method [22]. For `flare`, we set the tuning parameter to $\lambda = c\sqrt{n}\Phi^{-1}(1-\alpha/(2p))$ with $c = 1.1$ as suggested in Belloni et al. [2]. In an attempt to get as close as possible to the global optimum, we decrease the `prec` (precision) option to 1×10^{-10} and increase `max.ite` (maximum iterations) to 1×10^{7}.

5.3 Simulation results

We present results for four different measures of success based on different data analysis objectives. These are prediction risk, which evaluates how well we can predict a new Y given a new X; consistency, which measures how far the estimated $\widehat{\beta}$ is from β^*; F-score, which considers how well a method does at model selection; and risk estimation, which evaluates how close each estimate of the risk is to the true risk in a neighborhood of β^*.

In general, the unbiased risk-estimation methods we develop perform quite well across simulation conditions and evaluation metrics. The 2-stage method described in Section 4.3 also performs extremely well and warrants further investigation. Standard 10-fold cross-validation performs adequately while the behavior of scaled-sparse regression, $\sqrt{\text{lasso}}$ variants, and modified cross-validation depends strongly on the simulation condition. In particular, these modern methods often underperform the methods presented in this paper. Lastly, note that the refitted $\sqrt{\text{lasso}}$ essentially always improves on the performance of $\sqrt{\text{lasso}}$ and hence should strongly be considered in practice.

5.3.1 Prediction risk

Prediction risk is an important criterion as it is often a major goal in modern data analysis applications. For a new independent draw from $Z^\top = (X^\top, Y)$ from the same distribution that generates the data in equation (1), let $\mathbb{E}(Y - X^\top\beta)^2$ be the prediction risk (where $\mathbb{E}$ is the integration operator with respect to Z only, conditional on the training data). Then we define

$$\text{Risk}(\beta) = \mathbb{E}(Y - X^\top\beta)^2 - \sigma^2.$$

to be prediction risk ignoring the noise level. For these simulations, we approximate this quantity with the average squared error over 5000 test observations, but continue to denote it $\text{Risk}(\beta)$.

We present boxplots for the prediction risk of the selected models in Figure 4 and Figure 5 for SNR 0.1 and 10 respectively. Across all conditions, the 2-stage procedure performs consistently well, frequently resulting in the best, or nearly the best performance. Ten-fold cross-validation also performs quite well with the exception of the low-SNR, low-ρ scenario (Figure 4, columns 1 and 3), when its performance is merely average. R-CV-2, R-RCV-2, and R-CV-logn perform adequately in all settings with the caveat that R-CV-2 is better than R-CV-logn with high SNR and the opposite result for low SNR. MCV is relatively worse in general, but especially so with high SNR, while SQRT and refitted SQRT are quite bad with large ρ or large SNR. Finally, R-RMLE-2 is worse with low SNR, especially when the true model is sparse (Figure 4, top row), and SSR works poorly when the SNR is high and β^* is less sparse (Figure 5, bottom row).

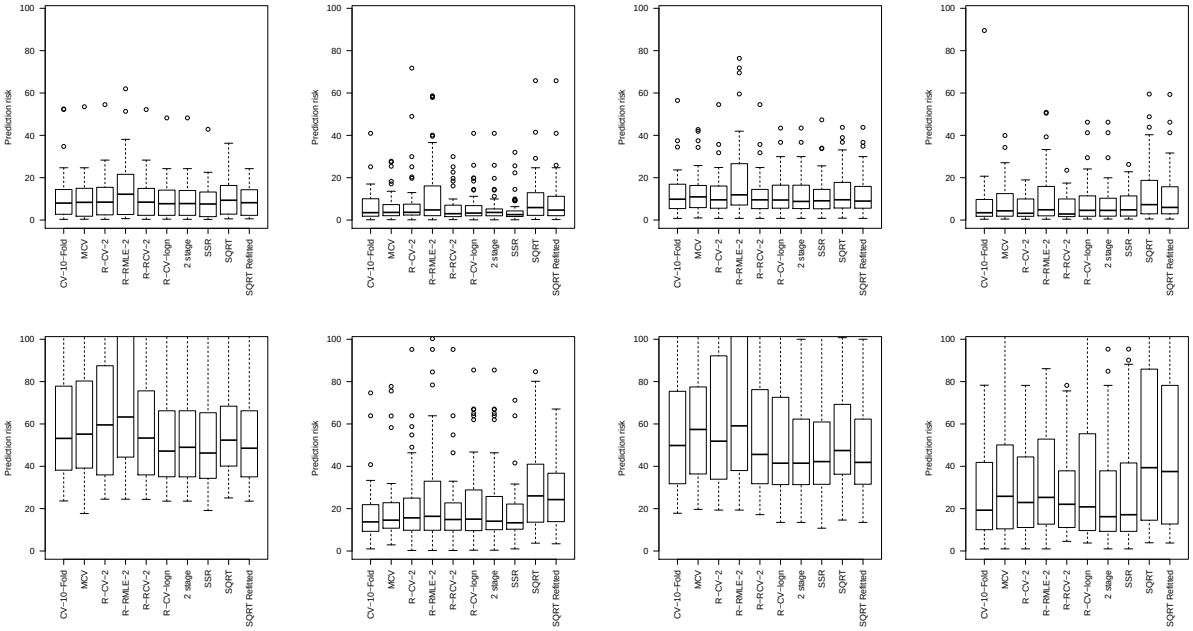

$p = 200, \rho = 0.1$ $p = 200, \rho = 0.8$ $p = 1000, \rho = 0.1$ $p = 1000, \rho = 0.8$

Figure 4: Comparison of prediction risk for $SNR = 0.1$. Top row: $\alpha = 0.4$. Bottom row: $\alpha = 0.7$.

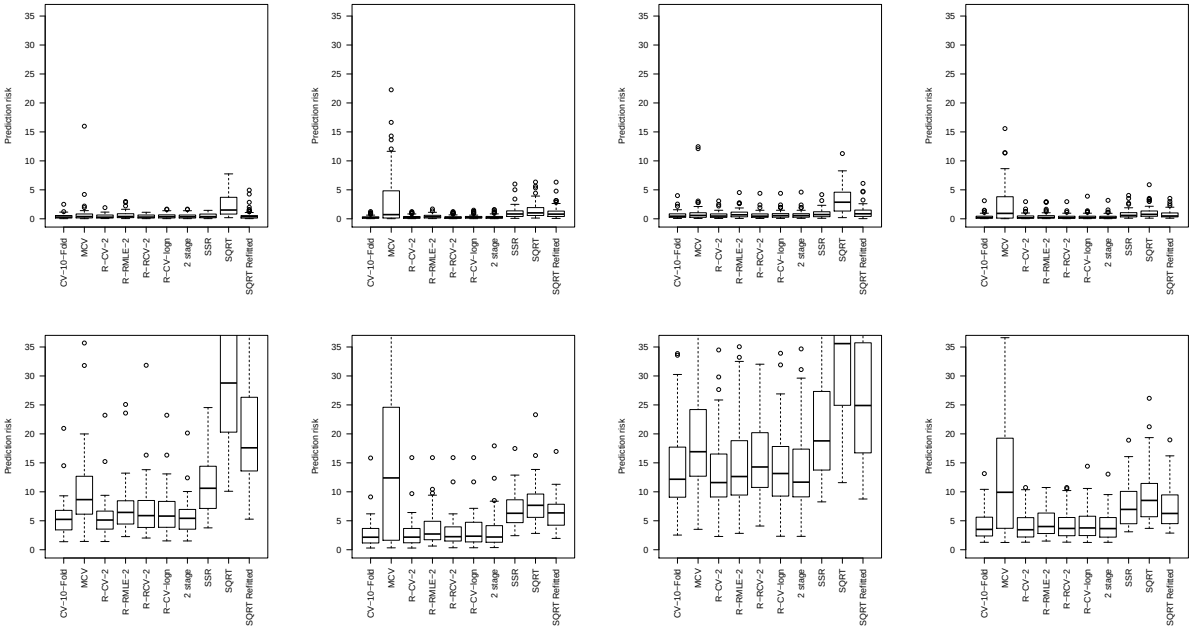

$p = 200, \rho = 0.1$ $p = 200, \rho = 0.8$ $p = 1000, \rho = 0.1$ $p = 1000, \rho = 0.8$

Figure 5: Comparison of prediction risk for $SNR = 10$. Top row: $\alpha = 0.4$. Bottom row: $\alpha = 0.7$.

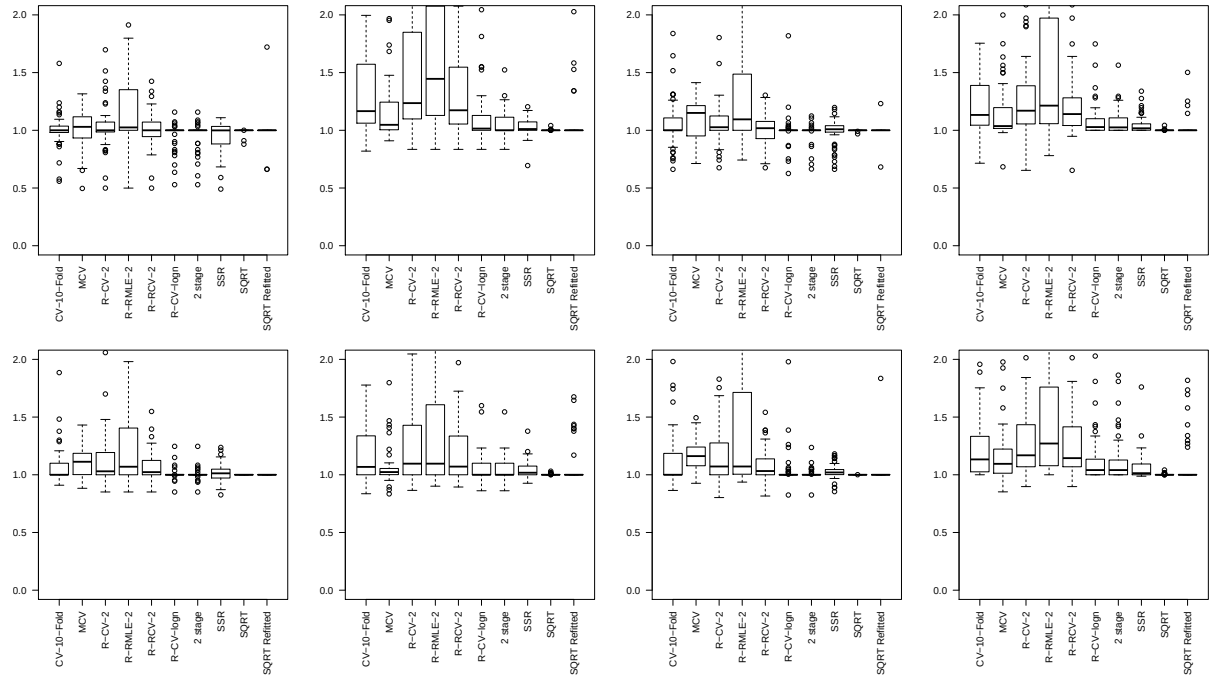

$p = 200, \rho = 0.1$ $p = 200, \rho = 0.8$ $p = 1000, \rho = 0.1$ $p = 1000, \rho = 0.8$

Figure 6: Comparison of consistency for $SNR = 0.1$. Top row: $\alpha = 0.4$. Bottom row: $\alpha = 0.7$.

5.3.2 Consistency

The second performance metric we use examines the ability of our procedures to produce accurate estimates of the true parameters. We examine a normalized version of the deviation between the estimated coefficients and the size of the parameter:

$$C(\widehat{\beta}) = \frac{\mathbb{E}\left|\left|\widehat{\beta} - \beta^*\right|\right|_2^2}{||\beta^*||_2^2}.$$

Thus, smaller values are better, and values near 1 often represent overly sparse solutions: $\widehat{\beta} \approx 0 \Rightarrow \mathbb{E}\left|\left|\widehat{\beta} - \beta^*\right|\right| \approx ||\beta^*||$.

In the low-SNR condition (Figure 6), no procedure performs particularly well, as one would expect. The two versions of $\sqrt{\text{lasso}}$ nearly always select $\widehat{\beta} = 0$ with occasional exceptions which tend to result in larger $C(\widehat{\beta})$, especially in the high-correlation scenario. R-CV-logn and the 2-stage procedure often do adequately, while R-RMLE-2 is the worst. MCV has low variance across simulations, but is worse on average when the correlation is low (columns 1 and 3). Ten-fold CV, R-CV-2 and R-RCV-2 fall somewhere in the middle.

With high SNR (Figure 7), SSR, and both versions of $\sqrt{\text{lasso}}$ are quite poor. MCV is also bad unless the model is sparse and the correlation is low (first and third plots on the top row). The remaining methods all perform well with 10-fold CV, R-CV-2, R-CV-logn, and the 2-stage procedure performing slightly better than R-RMLE-2 and R-RCV-2.

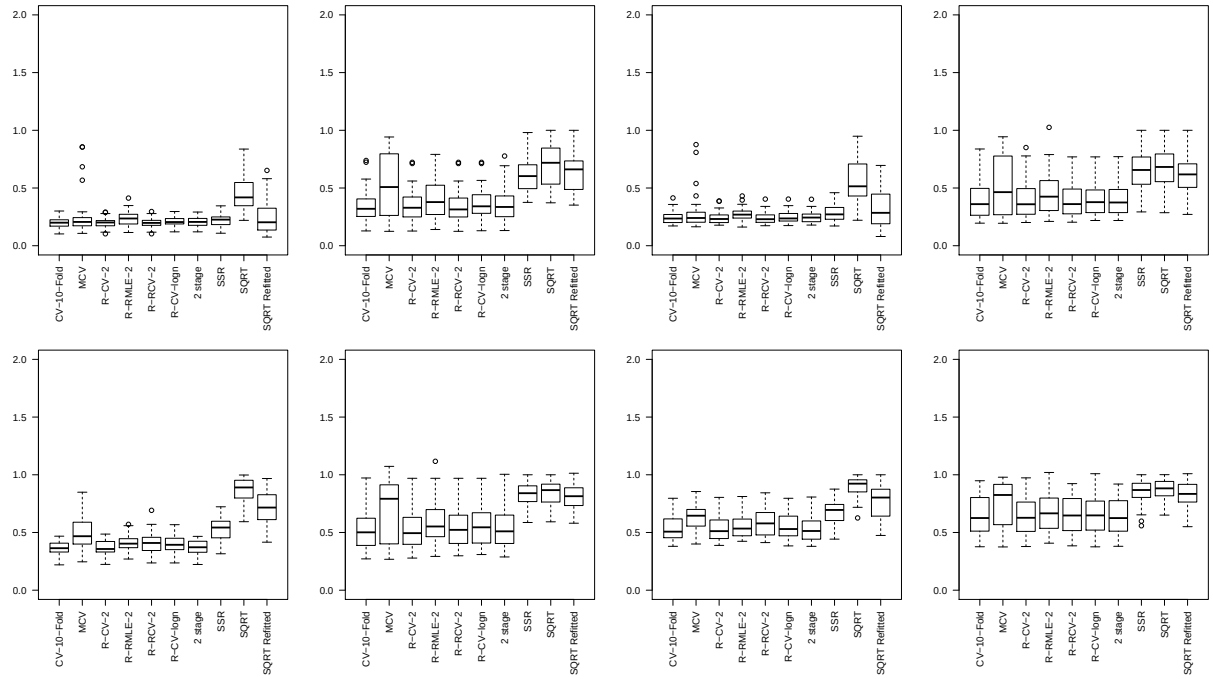

$p = 200, \rho = 0.1$ $p = 200, \rho = 0.8$ $p = 1000, \rho = 0.1$ $p = 1000, \rho = 0.8$

Figure 7: Comparison of consistency for $SNR = 10$. Top row: $\alpha = 0.4$. Bottom row: $\alpha = 0.7$.

5.3.3 F-score

To examine the ability of these procedures to perform model selection directly, we define the precision and recall for a particular β (recalling that $\mathcal{S} = \{j : |\beta_j| > 0\}$) to be

$$P(\mathcal{S}) = \frac{|\mathcal{S} \cap \mathcal{S}^*|}{|\mathcal{S}|} \qquad \text{and} \qquad R(\mathcal{S}) = \frac{|\mathcal{S} \cap \mathcal{S}^*|}{|\mathcal{S}^*|},$$

respectively. To parsimoniously represent both precision and recall at the same time, we use the F-score (sometimes referred to as the $F1$-score), which is the harmonic mean of the precision and recall:

$$F(\mathcal{S}) = \frac{2R(\mathcal{S})P(\mathcal{S})}{R(\mathcal{S}) + P(\mathcal{S})} = \frac{2}{\frac{1}{R(\mathcal{S})} + \frac{1}{P(\mathcal{S})}}.$$

Observe that $F(\mathcal{S})$ is equal to one if and only if $R(\mathcal{S})$ and $P(\mathcal{S})$ are both equal to one and equal to zero if either $R(\mathcal{S})$ or $P(\mathcal{S})$ are equal to zero. Thus, higher values represent better performance.

For the low SNR case (Figure 8), the 2-stage procedure, R-CV-logn, and both versions of $\sqrt{\text{lasso}}$ are quite poor (note that $\sqrt{\text{lasso}}$ and refitted $\sqrt{\text{lasso}}$ necessarily have the same F-score). If the design correlation, ρ, is low (columns 1 and 3), SSR is the best, though MCV also works quite well. The performance of MCV here depends much less on the correlation than in the previous sections, and it works well in high-ρ conditions as well, where SSR is not very good. Ten-fold CV, R-CV-2, R-RMLE-2, and R-RCV-2 have more variability than MCV, so while they occasionally perform better, they often perform worse, and the median

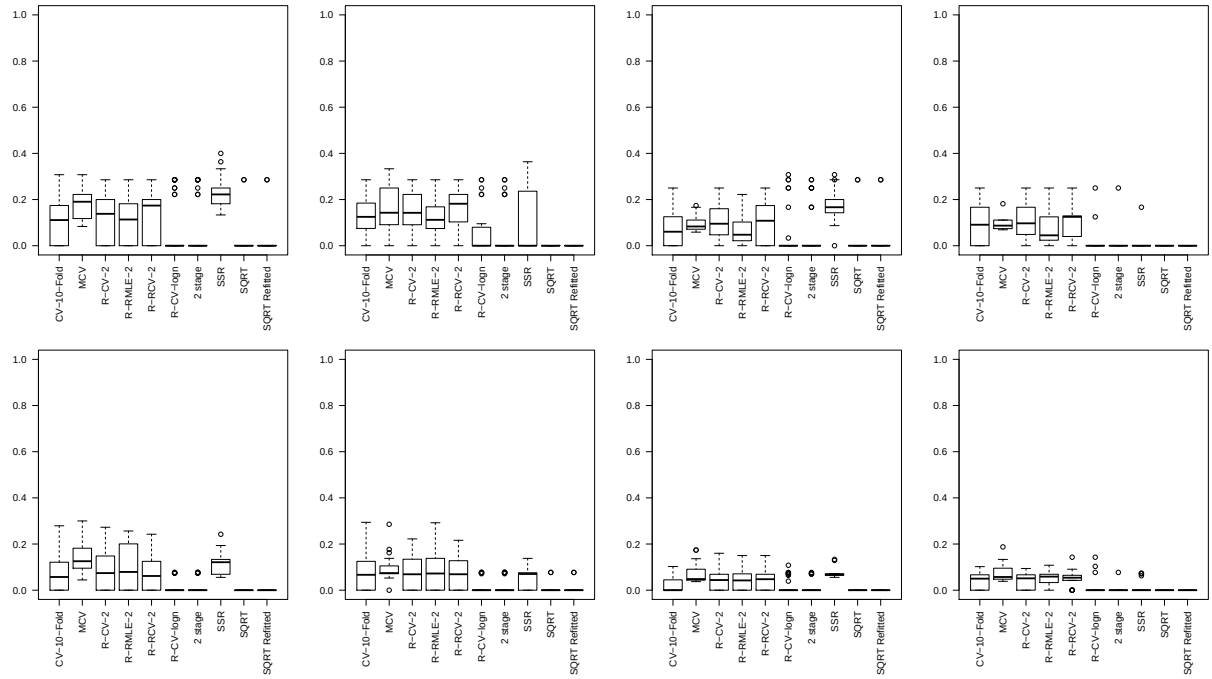

Figure 8: Comparison of F-score for $SNR = 0.1$. Top row: $\alpha = 0.4$. Bottom row: $\alpha = 0.7$.

performance tends to be slightly worse than MCV. Note that $\sqrt{\text{lasso}}$ is by far the worst performer in both cases as it tends to select an overly sparse model.

For the high SNR case (Figure 9), the 2-stage method works well across all settings of α and ρ as does R-CV-logn, though slightly less so. When α is small, then $\sqrt{\text{lasso}}$ has good F-score performance, but it is the worst when α is large. SSR works reasonably well except when α and ρ are both high. MCV is similar, though the behavior is less pronounced. R-RCV-2 works well for high α but is merely average if α is low. Ten-fold CV, R-CV-2, and R-RMLE-2 always fall in that order respectively and seem independent of the different scenarios, never performing particularly well relative to the best methods, or too poorly relative to the worst. As such, 10-fold CV strictly dominates the other two and R-RMLE-2 should be avoided.

5.3.4 Risk estimation

Sometimes, it becomes important to directly estimate the risk of a procedure to evaluate its performance with the same data that is used to train it. Here, we investigate the risk estimation property of both K-fold CV and $\widehat{R}(\widehat{\sigma}, 2/n)$ for a few choices of K and $\widehat{\sigma}^2$. As MCV, SSR, the 2-stage method, and $\sqrt{\text{lasso}}$ are really model selection/estimation procedures, we leave them out of these results. The goal here is to determine whether equation (3) can yield risk estimates in the high-dimensional setting the same way AIC can when $p < n$. Hence, we set $C_n = 2/n$ as this would be the unbiased choice if $\widehat{\sigma}^2$ were unbiased as well.

Using, for example, $\widehat{R}$ to both choose $\widehat{\lambda}$ and evaluate the risk $\widehat{\beta}(\widehat{\lambda})$ conflates $\widehat{R}$'s performance at tuning parameter selection and risk estimation. Hence, for this section, we use as

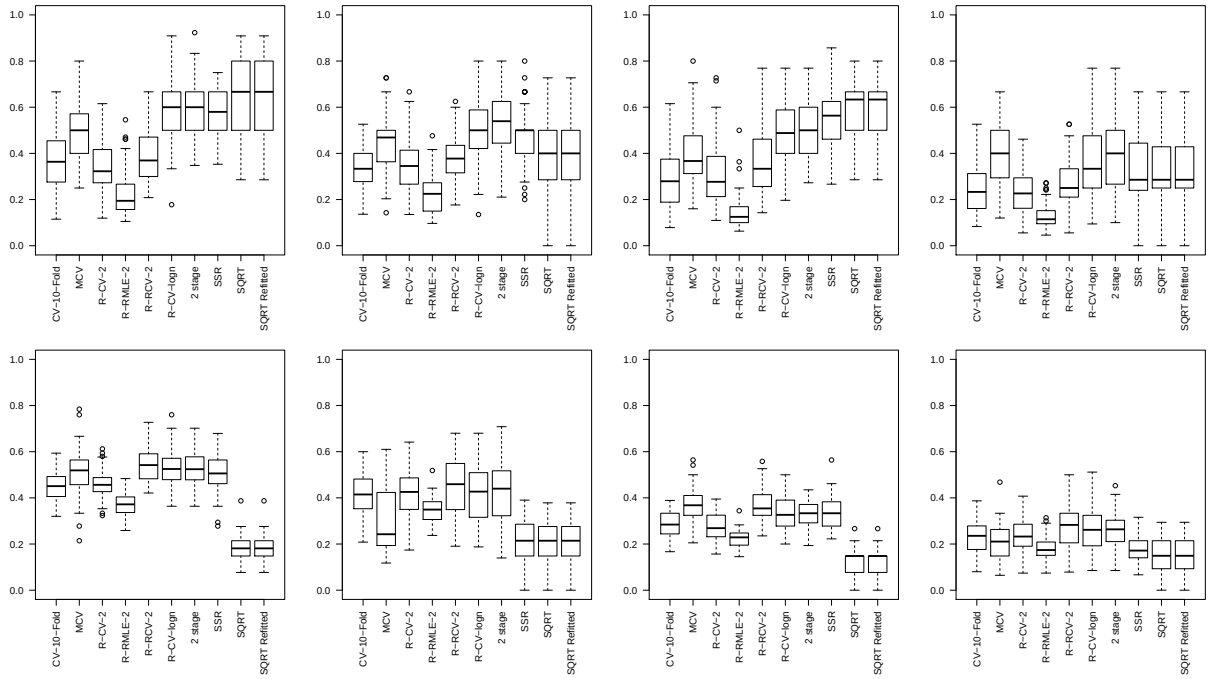

$p = 200, \rho = 0.0$ $p = 200, \rho = 0.8$ $p = 1000, \rho = 0.0$ $p = 1000, \rho = 0.8$

Figure 9: Comparison of F-score for $SNR = 10$. Top row: $\alpha = 0.4$. Bottom row: $\alpha = 0.7$.

a β^*-estimation procedure the oracle least squares estimator. That is, we set

$$\widehat{\beta}_O = \underset{\beta}{\operatorname{argmin}} ||Y - \mathbb{X}_{S^*}\beta||_2^2$$

and then calculate $\widehat{R}_{\widehat{\beta}_O}(\widehat{\sigma}^2_{CV}, 2/n)$, $\widehat{R}_{\widehat{\beta}_O}(\widehat{\sigma}^2_{RCV}, 2/n)$, and $\widehat{R}_{\widehat{\beta}_O}(\widehat{\sigma}^2_{RMLE}, 2/n)$ where $\widehat{\sigma}^2$ is estimated with the relevant high-dimensional variance estimator. We also include 2-Fold CV and 10-Fold CV (denoted CV-2-Fold and CV-10-Fold, respectively). This choice of β^* estimation procedure is still a function of the data, and hence is random, but should be in a neighborhood of β^*.

We find that for sparse models (Figure 10 and Figure 11, top rows), there is very little difference between these five procedures: all are unbiased on median, though 2-Fold CV has slightly larger variance. However, with less sparse models, 2-Fold CV greatly overestimates the risk, while 10-Fold CV is quite accurate. Table 2 presents the mean squared error of each method to summarize the figures. More concretely, Table 2 shows the squared difference between the risk estimate and the true risk (1 in all cases), averaged across the simulation runs—the risk of the risk estimator. Looking down the table for dense models, we see that for small SNR, R-RCV-2 is the best method although 10-Fold CV and R-CV-2 are close behind in terms of MSE. This is because the small negative bias of R-RCV-2 is outweighed by the smaller variance it has relative to 10-fold CV and R-CV-2, which are relatively unbiased. With high SNR, R-RCV-2 is terrible with high positive bias and huge variance, worse than even 2-Fold CV. Note that R-RCV-2 uses a version of 2-Fold CV to estimate σ^2. Here, 10-Fold CV is easily the best, R-CV-2 has low bias, but relatively large variance, while R-RMLE-2 has a pronounced downward bias with small variance.

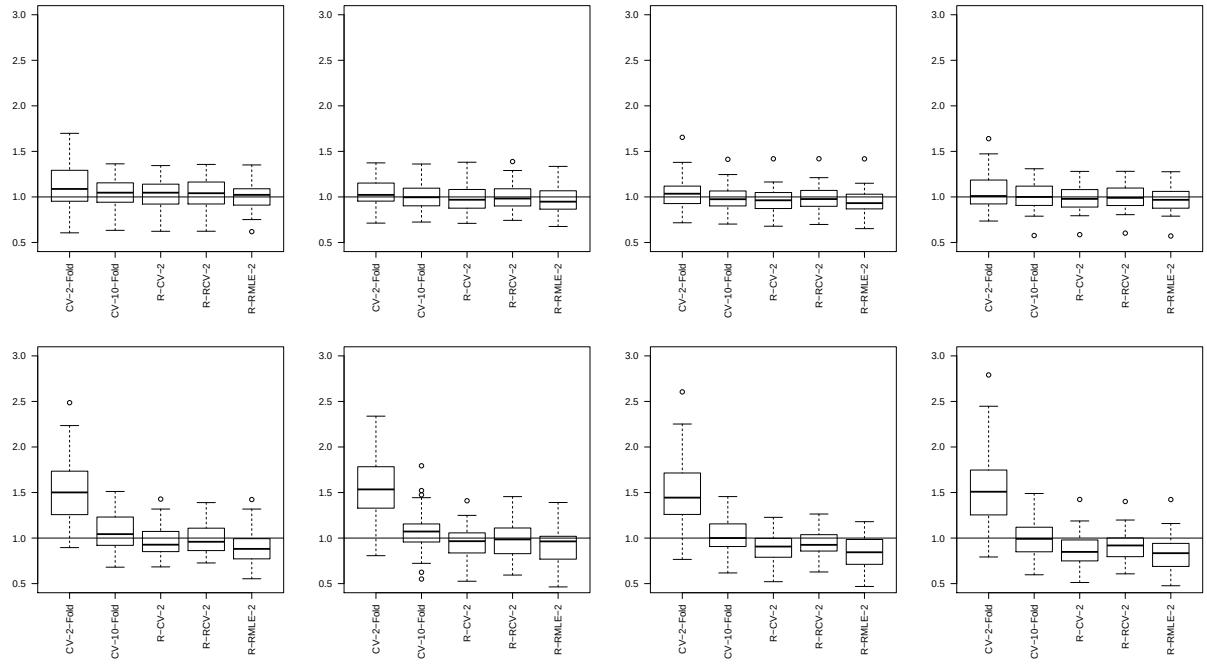

$p = 200, \rho = 0.1$ $p = 200, \rho = 0.8$ $p = 1000, \rho = 0.1$ $p = 1000, \rho = 0.8$

Figure 10: Comparison of risk estimation for $SNR = 0.1$. Top row: $\alpha = 0.4$. Bottom row: $\alpha = 0.7$.

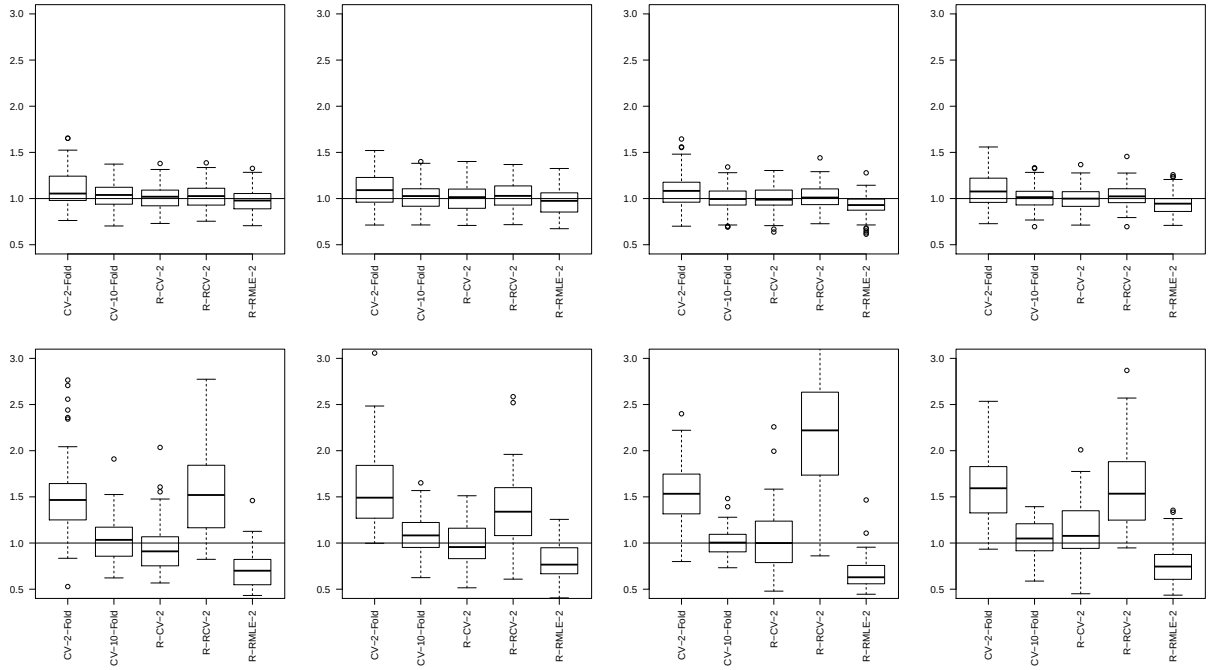

$p = 200, \rho = 0.1$ $p = 200, \rho = 0.8$ $p = 1000, \rho = 0.1$ $p = 1000, \rho = 0.8$

Figure 11: Comparison of risk estimation for $SNR = 10$. Top row: $\alpha = 0.4$. Bottom row: $\alpha = 0.7$.

SNR	α	p	ρ	CV-2-Fold	CV-10-Fold	R-CV-2	R-RCV-2	R-RMLE-2
0.1	0.4	200	0.1	0.060	**0.029**	**0.028**	**0.028**	**0.026**
0.1	0.4	200	0.8	0.027	**0.022**	**0.021**	**0.019**	**0.021**
0.1	0.4	1000	0.1	0.030	**0.019**	**0.018**	**0.018**	**0.022**
0.1	0.4	1000	0.8	0.041	**0.022**	**0.019**	**0.020**	**0.019**
0.1	0.7	200	0.1	0.414	0.047	**0.028**	**0.028**	0.040
0.1	0.7	200	0.8	0.424	0.059	**0.034**	**0.034**	**0.038**
0.1	0.7	1000	0.1	0.391	0.032	0.032	**0.022**	0.054
0.1	0.7	1000	0.8	0.446	0.042	0.046	**0.032**	0.059
10	0.4	200	0.1	0.054	**0.022**	**0.021**	**0.021**	**0.019**
10	0.4	200	0.8	0.042	**0.023**	**0.024**	**0.025**	**0.022**
10	0.4	1000	0.1	0.051	**0.021**	**0.020**	**0.020**	**0.025**
10	0.4	1000	0.8	0.043	**0.016**	**0.018**	**0.019**	**0.018**
10	0.7	200	0.1	0.499	**0.059**	0.085	0.720	0.119
10	0.7	200	0.8	0.533	**0.053**	**0.056**	0.365	0.074
10	0.7	1000	0.1	0.415	**0.024**	0.128	2.439	0.143
10	0.7	1000	0.8	0.488	**0.037**	0.117	0.696	0.101

Table 2: The mean squared error of all five risk estimators. Bold values indicate the best method(s) (those within .005 of the minimum) in each case.

6 Discussion

Substantial theory exists for the optimal choice of the tuning parameter for the lasso and related methods. These results, however, depend both on unknown properties of the data generating process and unknown constants. Though there are many data-dependent methods for choosing the tuning parameters, there is a distinct lack of guidance in the literature about which method to use. This uncertainty is even more pronounced when faced with high-dimensional data where $p \gg n$.

We give examples that show that one commonly advocated approach, a generalized information criterion which has desirable theoretical properties in low dimensions, would necessarily choose the saturated model with $\lambda = 0$ when $p > n$. Therefore, we propose a risk estimator motivated by an unbiased estimator of prediction risk. This estimator requires three ingredients: an estimate of the degrees of freedom ($\widehat{\mathrm{df}}$), a constant that may depend on n (C_n), and an estimator of the variance ($\widehat{\sigma}^2$). While the degrees of freedom for the lasso problem is well understood, the other two choices are much less so. In particular, high-dimensional variance estimation is a difficult problem in its own right.

In this paper, we investigate the performance of a number of proposed methods for selecting λ in high-dimensional lasso problems. Our results supplement and elaborate upon those in Flynn et al. [12] that apply to the low-dimensional setting ($p < n$). In general, the unbiased-risk-estimation methods we present perform consistently well across conditions. They exhibit many of the properties common to the AIC-vs.-BIC debate (BIC selects smaller models, AIC is better for prediction) as well as some variation across variance estimators due to estimation bias. We further propose a novel two-stage method in Section 4.3 that

also performs consistently well and warrants further theoretical investigations.

Acknowledgements

The authors gratefully acknowledge support from the National Science Foundation (DH, DMS–1407543; DJM, DMS–1407439) and the Institute of New Economic Thinking (DH and DJM, INO14-00020).

References

[1] Akaike, H. (1974), "A new look at the statistical model identification," *Automatic Control, IEEE Transactions on*, **19**(6), 716–723.

[2] Belloni, A., Chernozhukov, V., and Wang, L. (2011), "Square-root lasso: pivotal recovery of sparse signals via conic programming," *Biometrika*, **98**(4), 791–806.

[3] Bühlmann, P., and van de Geer, S. (2011), *Statistics for high-dimensional data: Methods, theory and applications*, Springer, New York.

[4] Bunea, F., Tsybakov, A., and Wegkamp, M. (2007), "Sparsity oracle inequalities for the lasso," *Electronic Journal of Statistics*, **1**, 169–194.

[5] Chen, J., and Chen, Z. (2012), "Extended bic for small-n-large-p sparse glm," *Statistica Sinica*, **22**(2), 555.

[6] Efron, B. (1986), "How biased is the apparent error rate of a prediction rule?" *Journal of the American Statistical Association*, **81**(394), 461–470.

[7] Efron, B., Hastie, T., Johnstone, I., and Tibshirani, R. (2004), "Least angle regression," *The Annals of statistics*, **32**(2), 407–499.

[8] Fan, J., and Li, R. (2001), "Variable selection via nonconcave penalized likelihood and its oracle properties," *Journal of the American statistical Association*, **96**(456), 1348–1360.

[9] Fan, J., Guo, S., and Hao, N. (2012), "Variance estimation using refitted cross-validation in ultrahigh dimensional regression," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **74**(1), 37–65.

[10] Fan, Y., and Tang, C. Y. (2013), "Tuning parameter selection in high dimensional penalized likelihood," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **75**(3), 531–552.

[11] Feng, Y., and Yu, Y. (2013), "Consistent cross-validation for tuning parameter selection in high-dimensional variable selection," .

[12] Flynn, C. J., Hurvich, C. M., and Simonoff, J. S. (2013), "Efficiency for regularization parameter selection in penalized likelihood estimation of misspecified models," *Journal of the American Statistical Association*, **108**, 1031–1043.

[13] Friedman, J., Hastie, T., and Tibshirani, R. (2010), "Regularization paths for generalized linear models via coordinate descent," *Journal of statistical software*, **33**(1), 1.

[14] Hastie, T., Tibshirani, R., and Friedman, J. (2009), *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Springer Verlag.

[15] Hastie, T., Tibshirani, R., and Wainwright, M. J. (2015), *Statistical Learning with Sparsity: the Lasso and Generalizations*, CRC Press, Boca Raton, FL.

[16] Kim, Y., Kwon, S., and Choi, H. (2012), "Consistent model selection criteria on high dimensions," *The Journal of Machine Learning Research*, **13**(1), 1037–1057.

[17] Mallows, C. L. (1973), "Some comments on c_p," *Technometrics*, **15**(4), 661–675.

[18] Reid, S., Tibshirani, R., and Friedman, J. (2016), "A study of error variance estimation in lasso regression," *Statistica Sinica*, **26**(35–67).

[19] Schwarz, G. (1978), "Estimating the dimension of a model," *The Annals of Statistics*, **6**, 461–464.

[20] Stein, C. M. (1981), "Estimation of the mean of a multivariate normal distribution," *The Annals of Statistics*, **9**(6), 1135–1151.

[21] Sun, T., and Zhang, C.-H. (2012), "Scaled sparse linear regression," *Biometrika*, **99**(4), 879–898.

[22] Sun, T., and Zhang, C.-H. (2013), "Sparse matrix inversion with scaled lasso," *Journal of Machine Learning Research*, **14**, 3385–3418.

[23] Tibshirani, R. (1996), "Regression shrinkage and selection via the lasso," *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, **58**(1), 267–288.

[24] Tibshirani, R. J. (2013), "The lasso problem and uniqueness," *Electronic Journal of Statistics*, **7**, 1456–1490.

[25] Tibshirani, R. J., and Taylor, J. (2012), "Degrees of freedom in lasso problems," *Annals of Statistics*, **40**, 1198–1232.

[26] Wang, H., Li, R., and Tsai, C.-L. (2007), "Tuning parameter selectors for the smoothly clipped absolute deviation method," *Biometrika*, **94**(3), 553–568.

[27] Wang, H., Li, B., and Leng, C. (2009), "Shrinkage tuning parameter selection with a diverging number of parameters," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **71**(3), 671–683.

[28] Yu, Y., and Feng, Y. (2014), "Modified cross-validation for penalized high-dimensional linear regression models," *Journal of Computational and Graphical Statistics*, **23**(4), 1009–1027.

[29] Yuan, M., and Lin, Y. (2006), "Model selection and estimation in regression with grouped variables," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **68**(1), 49–67.

[30] Zhang, Y., and Shen, X. (2010), "Model selection procedure for high-dimensional data," *Statistical Analysis and Data Mining: The ASA Data Science Journal*, **3**(5), 350–358.

[31] Zhang, Y., Li, R., and Tsai, C.-L. (2010), "Regularization parameter selections via generalized information criterion," *Journal of the American Statistical Association*, **105**(489), 312–323.

[32] Zou, H., Hastie, T., and Tibshirani, R. (2007), "On the degrees of freedom of the lasso," *The Annals of Statistics*, **35**(5), 2173–2192.